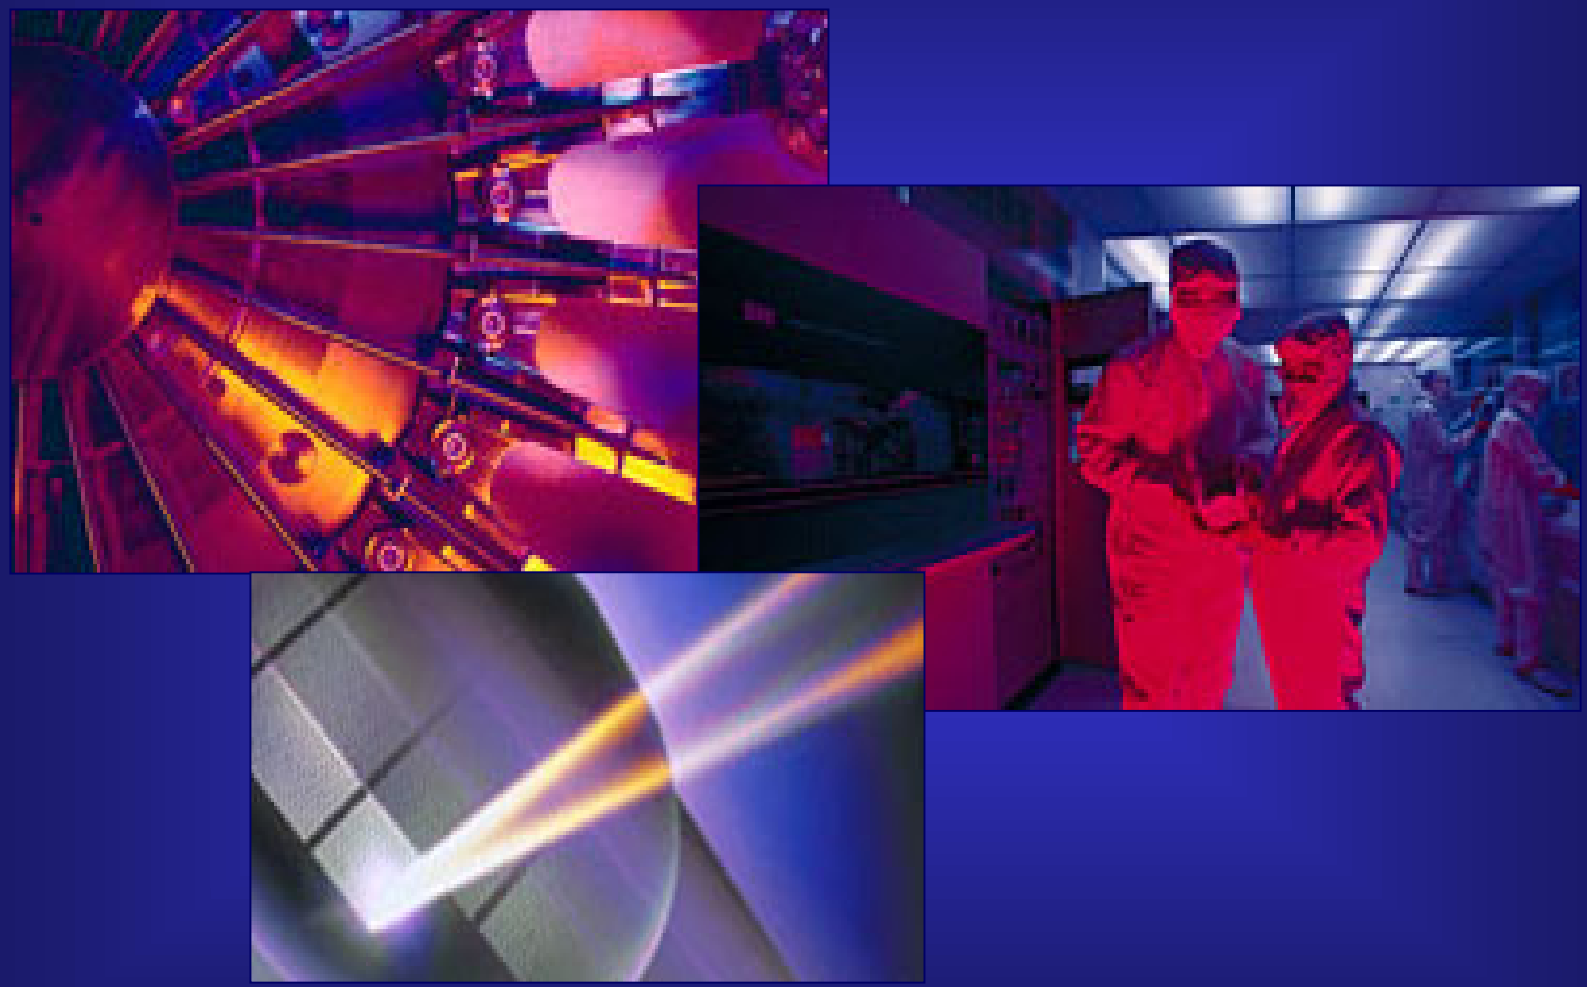

Daresbury MEW

December 9th 2004

Toby Smith
Solution Specialist
EMEA High Performance Computing

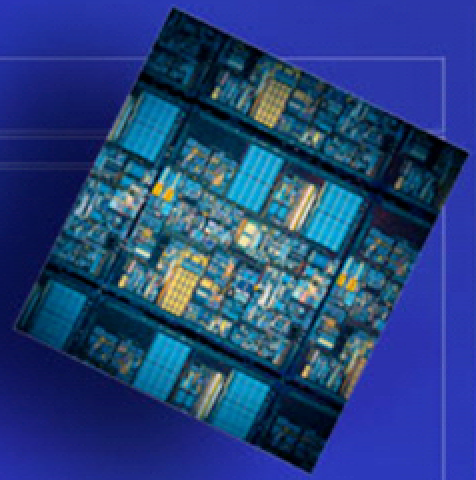

Intel Delivers HPC

intel

NASA Project Columbia – Out of this world Super Computing

10,240 Itanium processors (20 x 512P)

Climate Predictions **Weather Phenomena** **Re-entry Analysis**

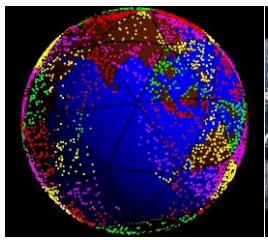

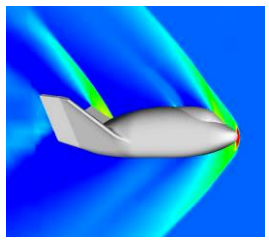

Deployed in record-breaking time

Built, installed, and operational on NASA Ames site in < 6 months, ~10X faster over proprietary supercomputers of its size

"The system will be built and integrated over the next 15 weeks (or 3 months). The first two nodes, in fact, were integrated June 28-30 and became operational in early July. This initial build doubled the current capacity at NASA Ames."
- NASA Press Release July 27, 2004

Delivering Record breaking computers In record breaking time

intel

*Other br

© Copy

Silicon Process Roadmap

Process Name	P856	P858	Px60	P1262	P1264	P1266	P1268
1st Production	1997	1999	2001	2003	2005	2007	2009
Lithography	250nm	180nm	130nm	90nm	65nm	45nm	32nm
Gate Length	200nm	130nm	<70nm	<50nm	<35m	<25nm	<18nm
Wafer Size	200mm	200mm	200/300 mm	300mm	300mm	300mm	300mm

Manufacturing

Development

Research

Extending Moore's Law

- New Intel technology generation every 2 years
- Intel R&D technologies drive this pace of innovation into the next decade

© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Multi-Core Transition

- Future products in development to deliver increasing performance per socket
- Scalable performance growth enabled through increased processor parallelism
- Most server applications are ready to take advantage of scalable architectures
- Intel delivering industry tools for multi-threaded software development

	Today Single Core	2005-2006 Dual Core	Future Multi-Core
Intel Inside Itanium 2	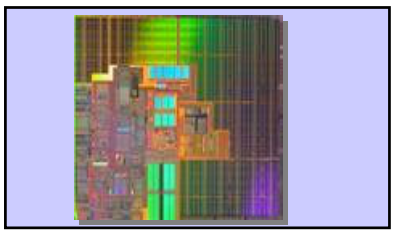	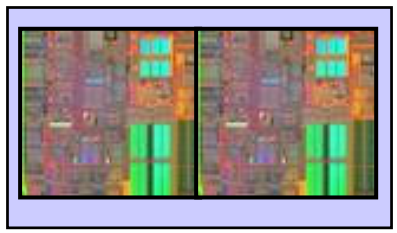	4 or more cores + Cache 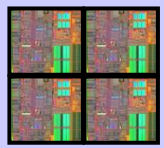
Intel Inside Xeon	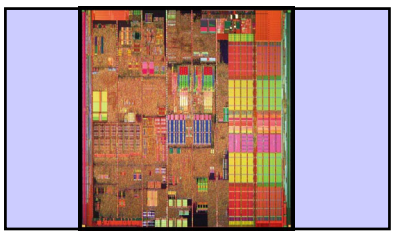	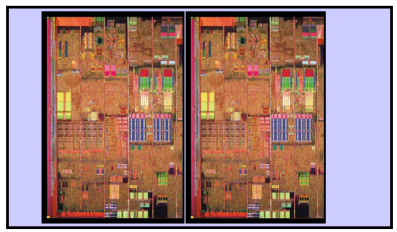	2 or more cores + Cache 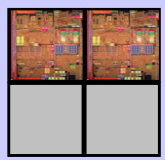

Dual Core Is Just The Beginning…

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Designing for Power

- Leakage current increases exponentially as process size decreases linearly
- Business as usual is not an option – 45nm CPU might need 1kW

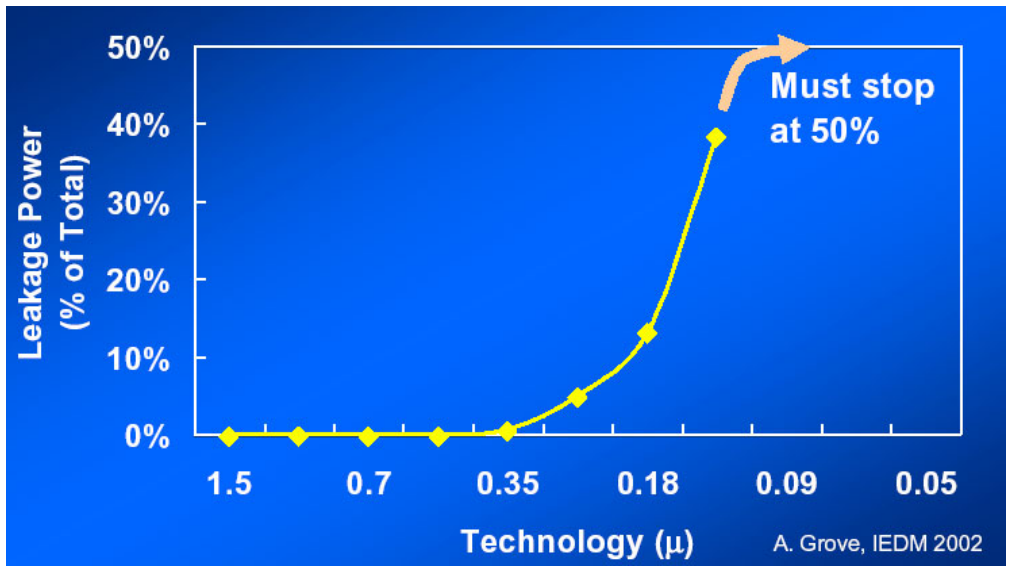

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Leakage Control

Body Bias

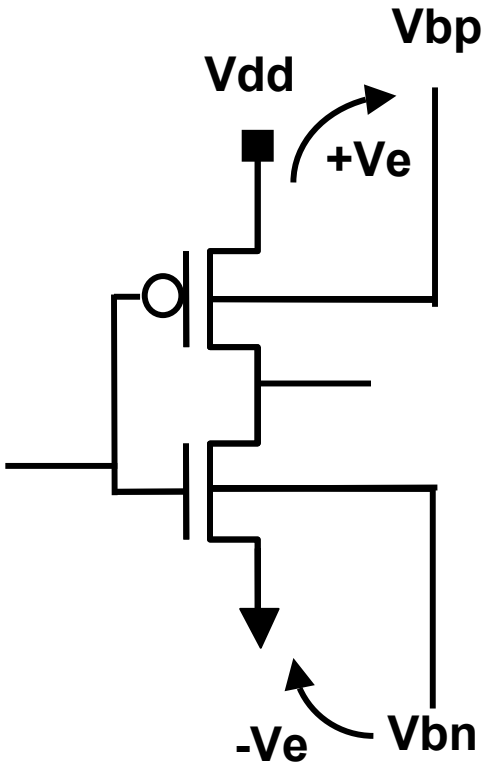

2-10X
Reduction

Stack Effect

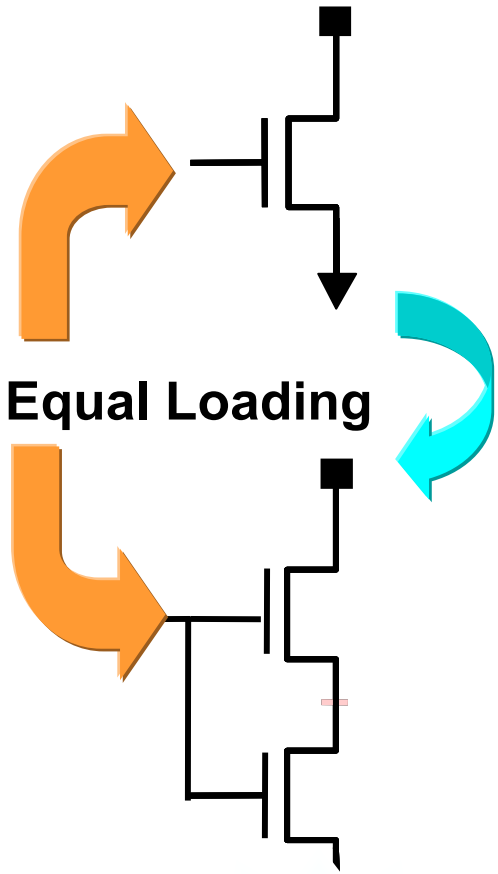

5-10X
Reduction

Sleep Transistor

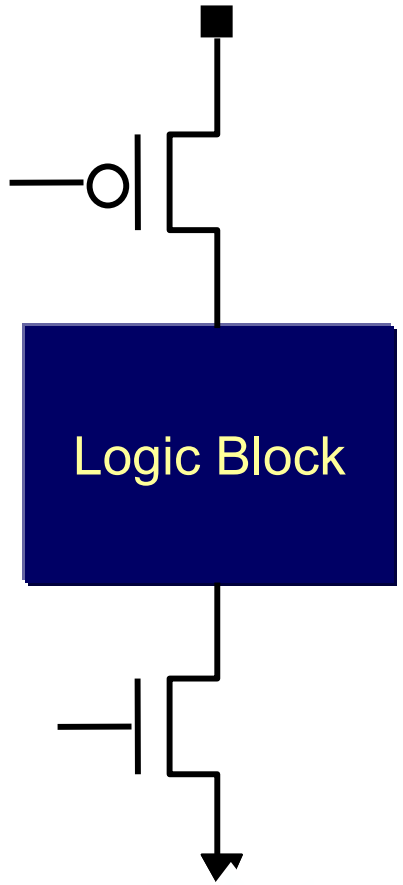

2-1000X
Reduction

*O[illegible]mes are the property of their respective ov[illegible]
© Copy[illegible]nt 2002-2003 Intel Corporation. All Rights Reserved.

Increase On-Die Memory

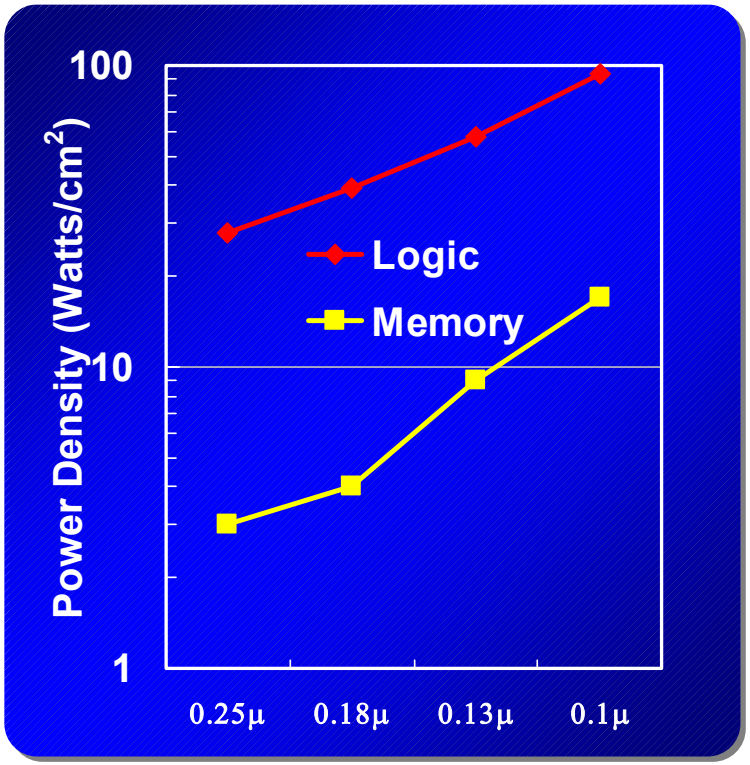

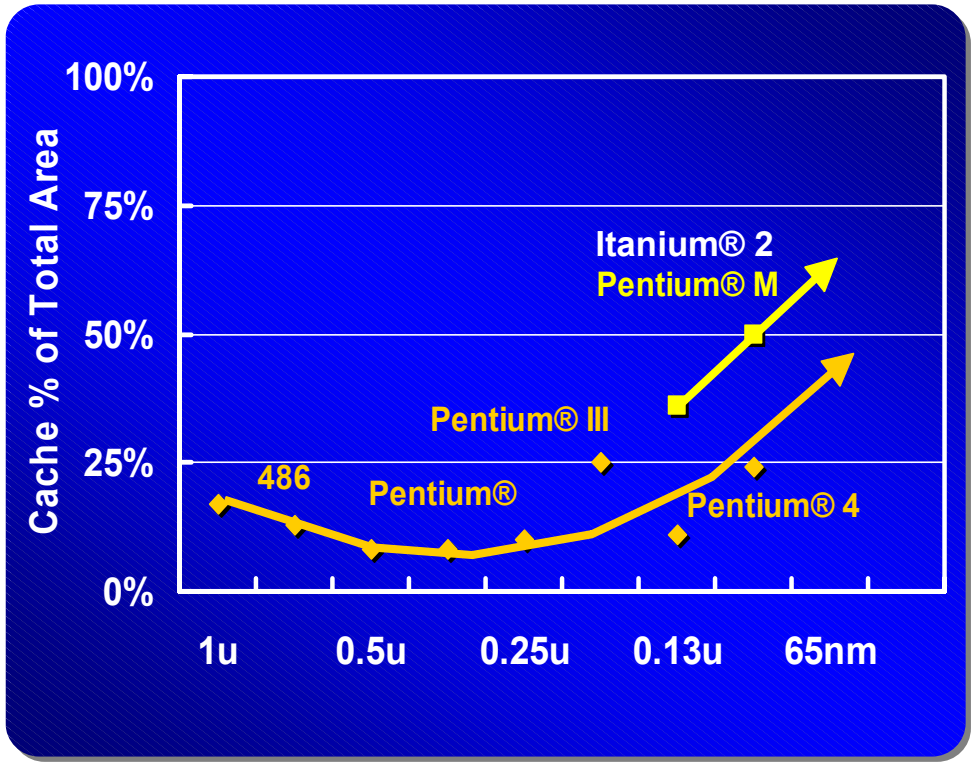

Large on-die memory provides:

- Increased Data Bandwidth & Reduced Latency
- Hence, higher performance for much lower power

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Chip Multi-Processing

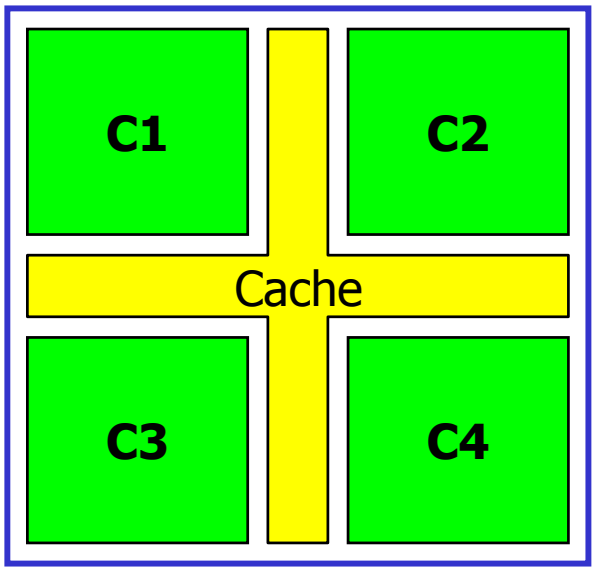

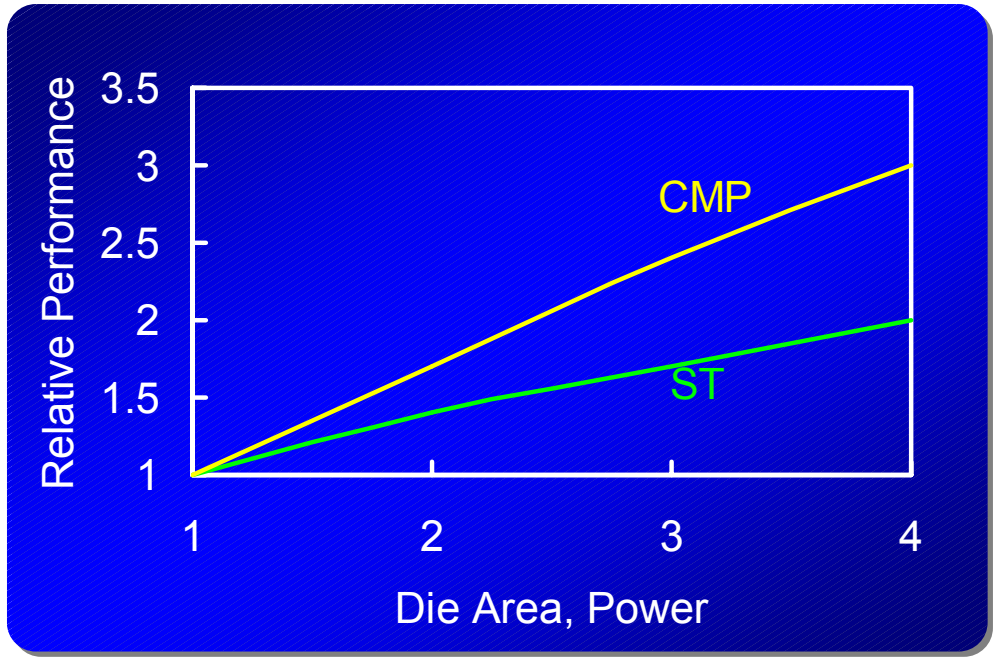

- Multi-core, each core Multi-threaded
- Shared cache and front side bus
- Each could have different Vdd & Freq
- Core hopping to spread hot spots

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Designing for Power

Special Purpose HW Engine

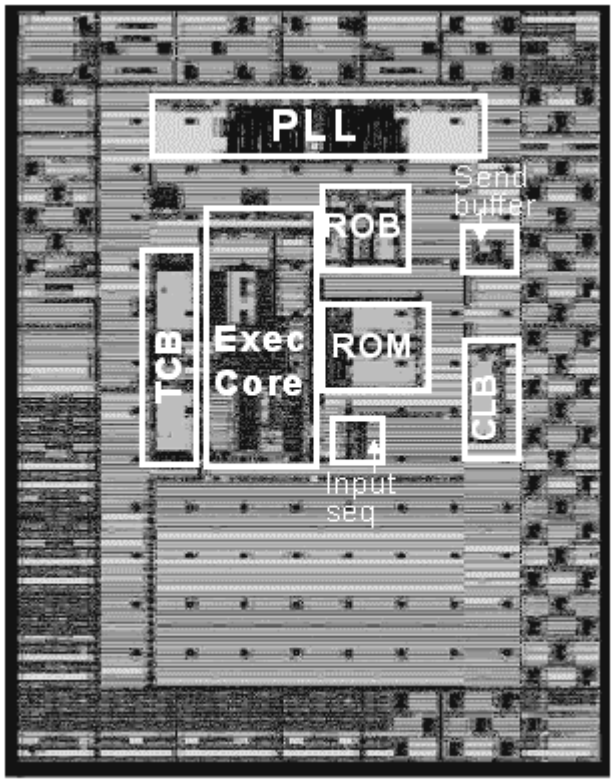

2.23 mm X 3.54 mm, 260K transistors

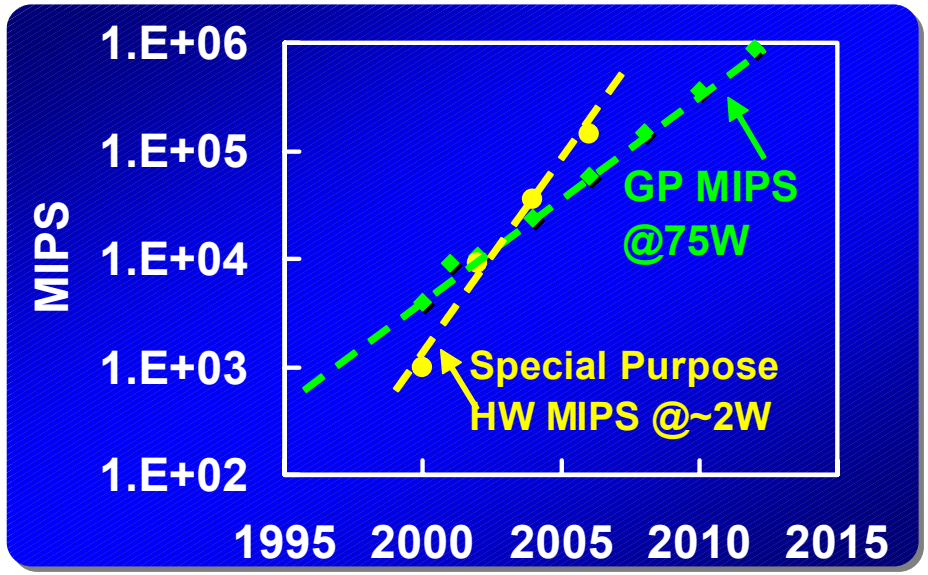

Opportunities:

- **Network processing engines**
- **MPEG Encode/Decode engines**
- **Speech engines**

Special purpose hardware = Best MIPS/Watt

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Intel® Server/WS Processor Roadmap

2003	2004	2005	Future
Intel® Itanium® 2 Processor (MP)	Madison9M**	Montecito** (Dual-Core)	Tukwila** (Multi-Core)
Intel® Itanium® 2 Processor (DP)	Fanwood** (DP)	Millington**	Dimona**
Intel® Itanium® 2 LV Processor (DP)	LV Fanwood**	LV Millington**	LV Dimona**

2003	2004	2005	Future
Intel® Xeon™ Processor MP	Gallatin**4M	Cranford** Potomac**	Tulsa** (Dual-Core)
Intel® Xeon™ Processor DP	Nocona**	Irwindale**	Future DP

All dates specified are target dates, are provided for planning purposes only and are subject to change.

**codename

intel

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Evolving 64-bit Computing

Current architecture or solutions	Transition benefits	Architecture of choice
RISC architecture	Exceptional performance with choice of OS, SW and HW vendors TODAY	intel inside ITANIUM 2 Mission critical 64-bit architecture. Premier performance, reliability and scalability; cost effective vs. RISC
IA-32 architecture	64-bit support via Intel® EM64T, great performance for 32-bit apps	intel inside XEON Mainstream 64-bit architecture; price/ performance and reliability

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Intel® Architecture in the Data Centre

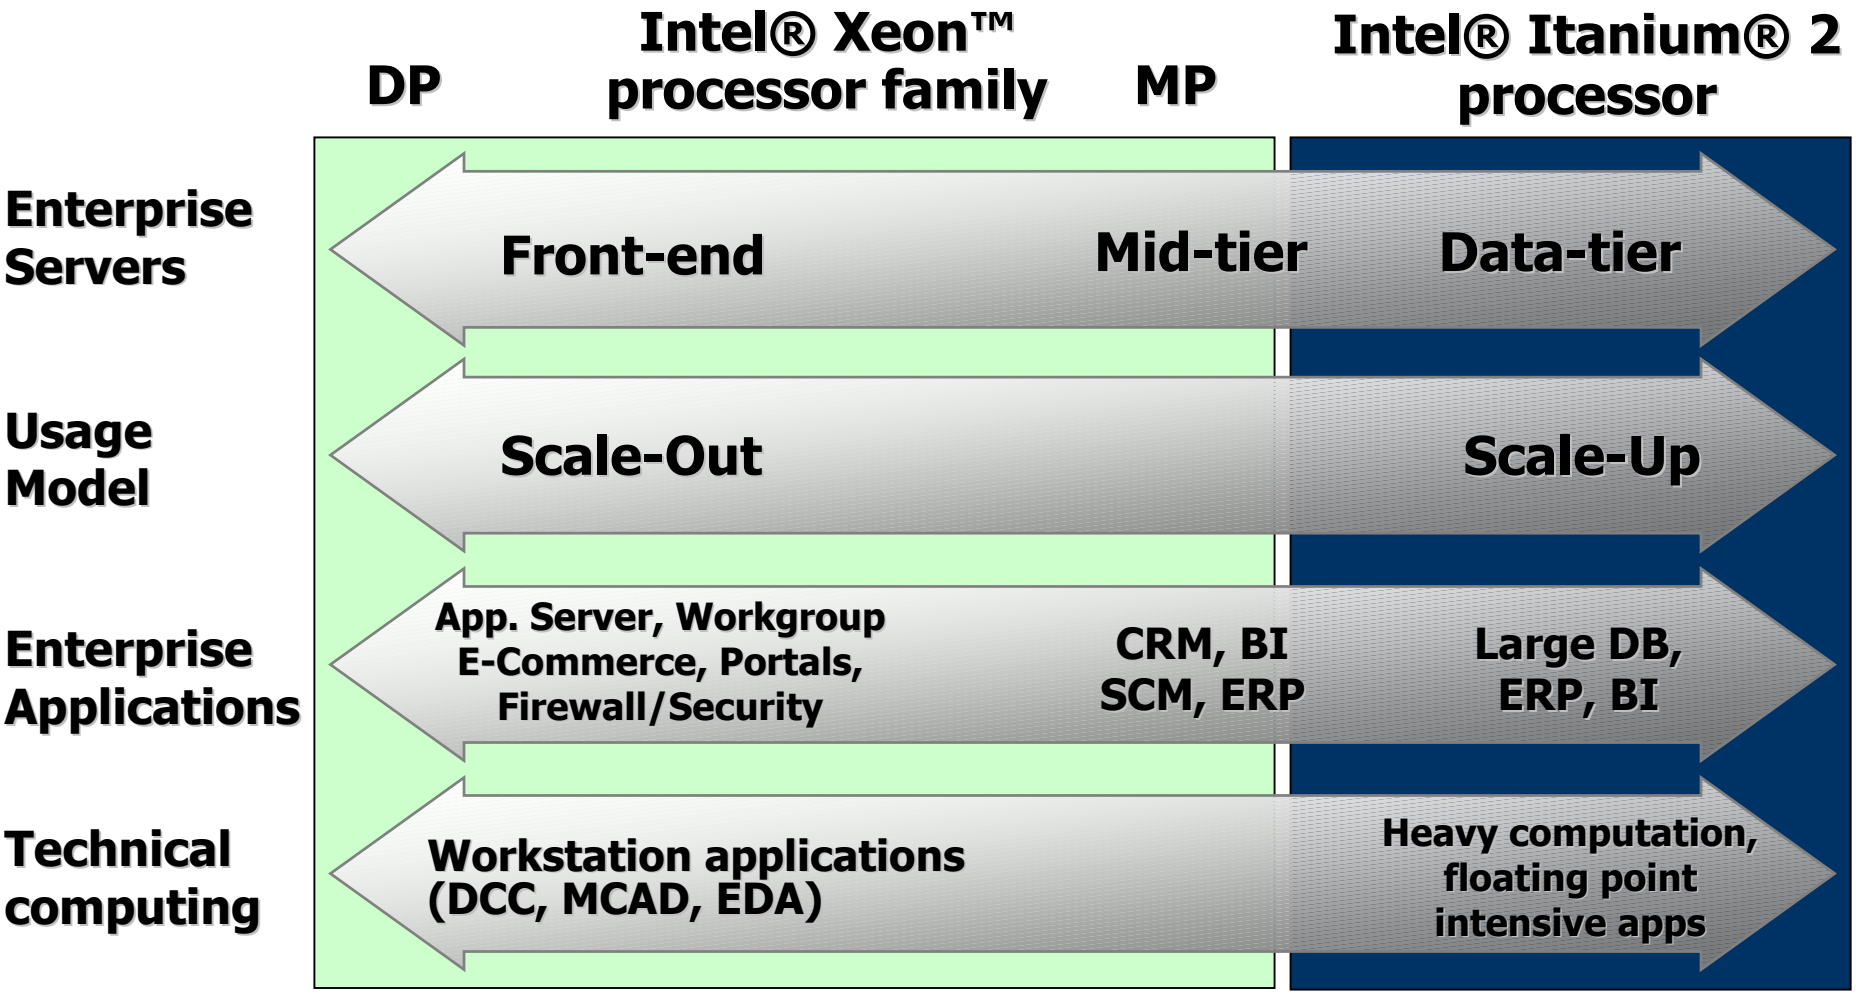

Two complementary 64-bit architectures to address different server market needs

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

IA-32 Processor Family

Multi-Processor (MP) Capable

Xeon™ MP Processor (Gallatin**) 3.0GHz, 4M L3

Cranford** =3.6 Potomac** =3.5GHz, 8MB L3

Tulsa** Dual Core

Whitefield**

Twin Castle** Chipset
PCI Express
EM64T
DDR2 Memory

Common Platform Architecture

Dual-Processor (DP) Capable

Nocona** 3.6GHz, 1MB L2

Irwindale** =3.6GHz, 2MB L2

Future DP Dual Core

Next Gen DP

Lindenhurst** Chipset
PCI Express
EM64T
DDR2 Memory

New DP Chipset
Enhanced Architecture
Virtualisation

Common Platform Architecture

2004 **2005** **2006+** **Next Gen**

** Codename

All products, dates and information are preliminary and subject to change without notice

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Streaming SIMD Extensions 3

- **SSE3 - 11 new instructions to allow compiler to boost application performance**

Floating Point to Integer	→	FISTTP
Complex Arithmetic	→	ADDSUBPD, ADDSUBPS, MOVDDUP, MOVSHDUP, MOVSLDUP
Video Encoding	→	LDDQU
SIMD Floating Point (AoS)	→	HADDPD, HSUBPD, HADDPS, HSUBPS
Thread Synchronization	→	MONITOR, MWAIT

http://www.intel.com/technology/itj/2004/volume08issue01/art02_compilers/vol8iss1_art02.pdf

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Intel® E7520 Chipset (Lindenhurst**)

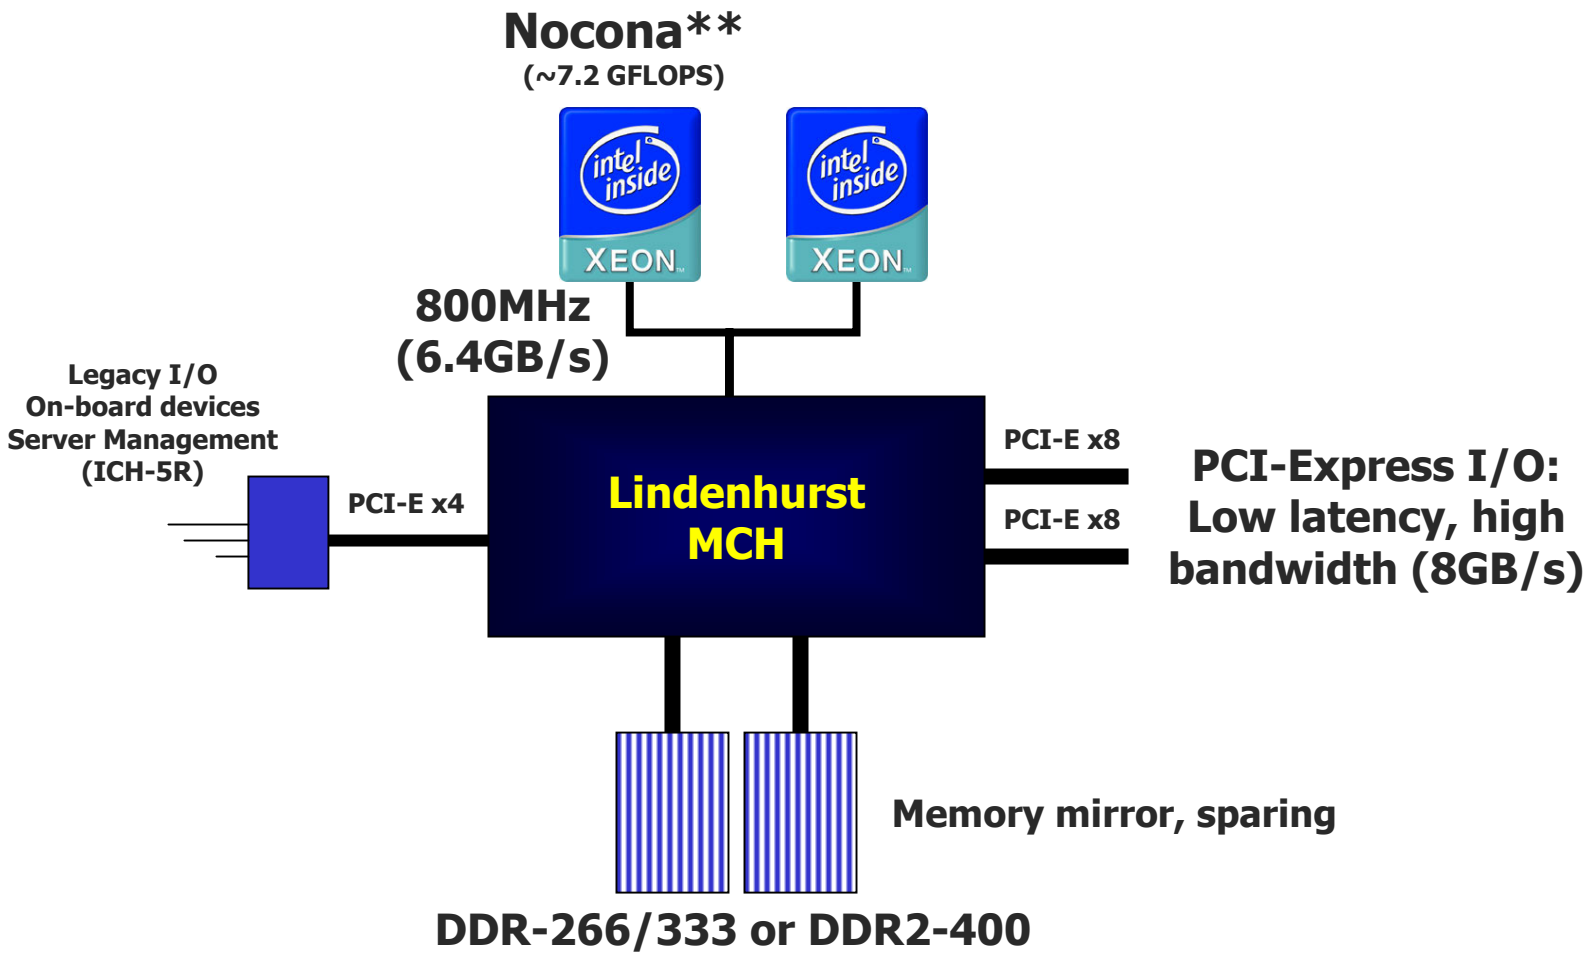

All products, features, dates, and figures are preliminary, for planning purposes only and are subject to change without notice.

**codename

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

64-Bit Extension Technology

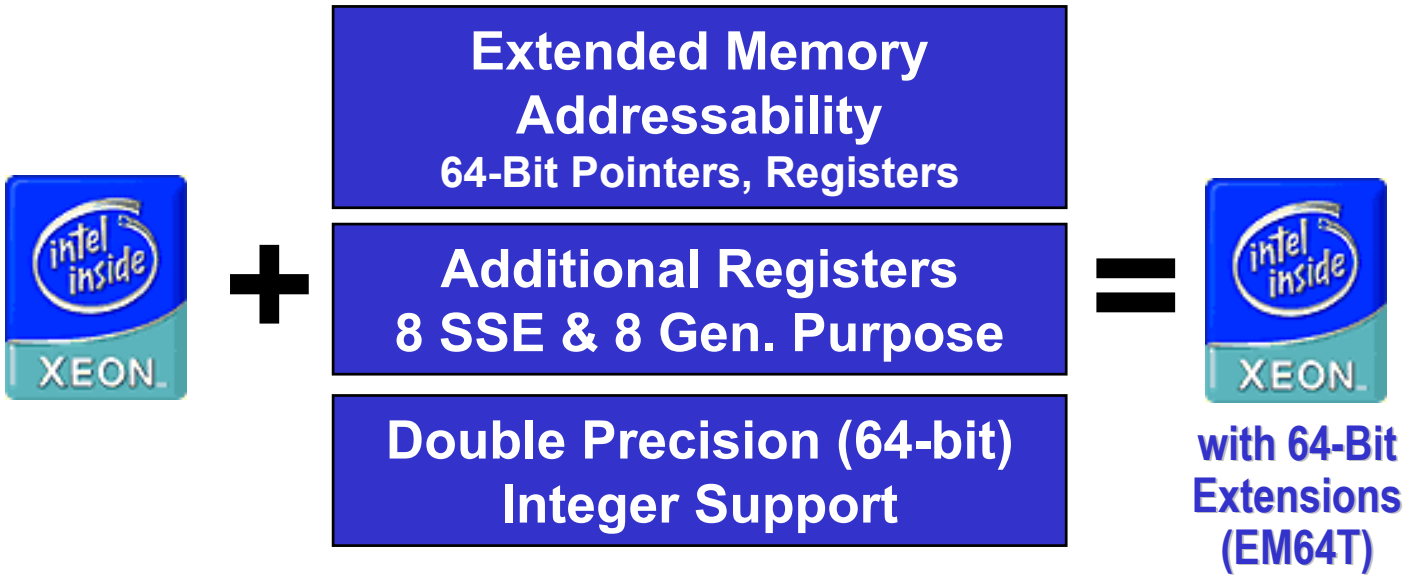

Application Support Modes

Mode	O/S	Application	Uses 64-bit Extensions?
Legacy Mode	32bit	16/32bit	No
Compatibility Mode	64bit	32bit	No
64-bit Mode	64bit	64bit	Yes

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Blade Clusters – The Next Wave?

- Compute blade
  - 2-way Intel® Xeon™ processor blade (1 slot)
  - 4-way Xeon processor blade (2 slots)
- Dual, redundant connectivity to backplane for LAN and storage networks
- 1 Gb Ethernet
- 2 GB Fibre Channel
- Future
  - Infiniband HBA/Switch
  - 10Gb Ethernet
  - Nocona/Potomac blades
- 7U Chassis
  - 28” Depth
  - 14 Compute Blade Slots
  - Dual-Dual Star Backplane Topology
  - Dual, Redundant Power Supply
  - Dual, Redundant Blowers

Enabling 1 TFLOP per Rack.

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Blade Power Consumption

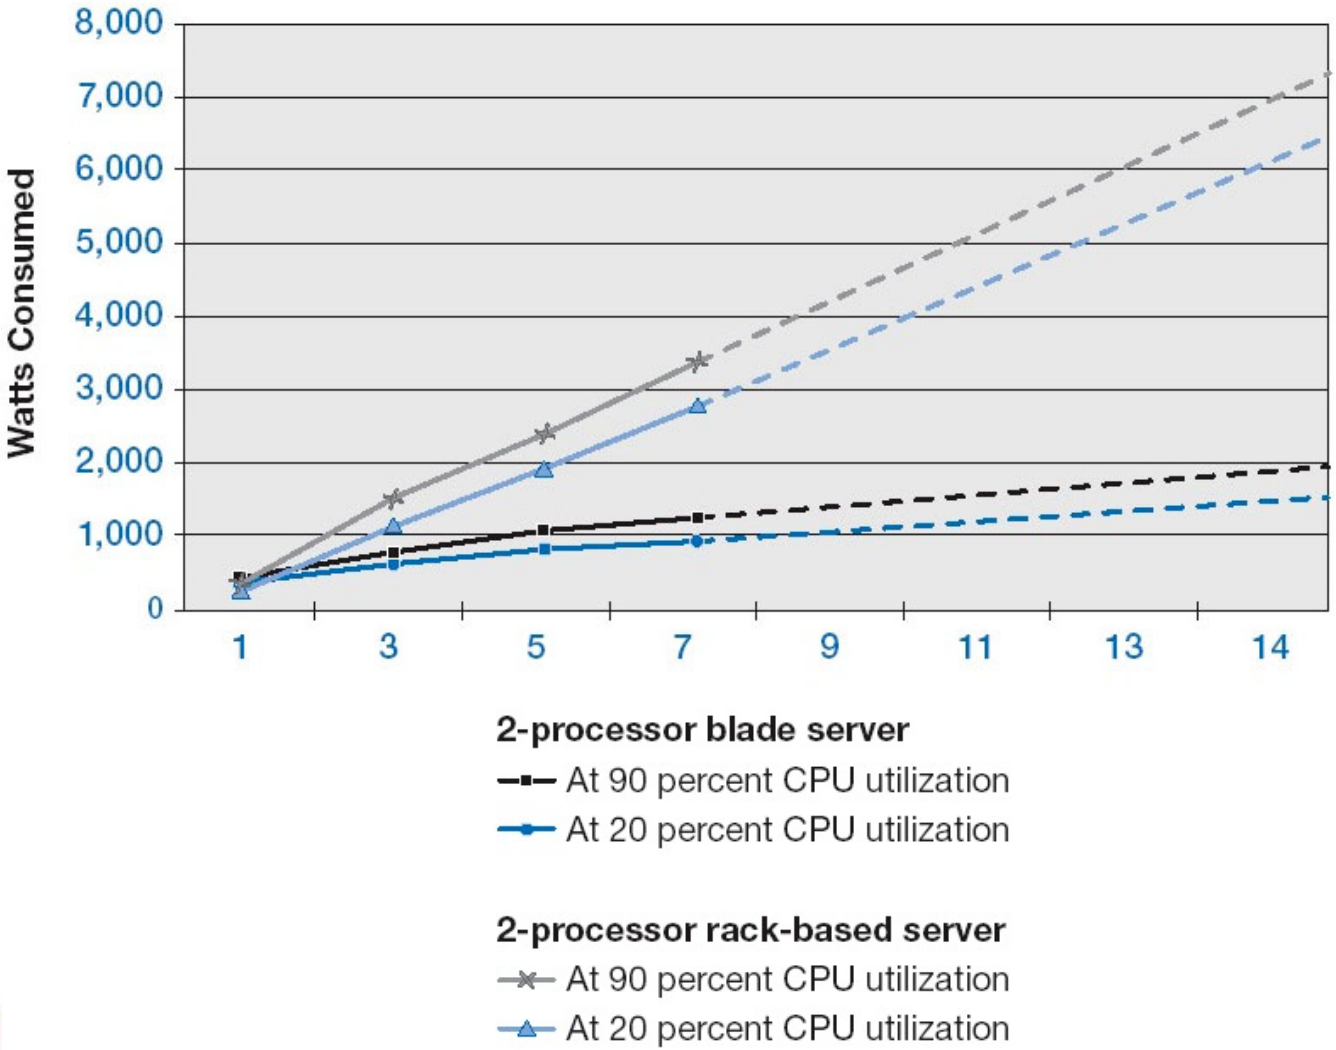

Source: IT@Intel White Paper, May 2004

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Introducing PCI Express™

- Why change?
  - Traditional I/O interconnects at practical limits
  - Lack of scaling with frequency and voltage
  - Gap between processing power and I/O bandwidth
- PCI Express™ will deliver the bandwidth for a wide range of complimentary I/O technologies
  - Graphics, Serial ATA, 1/10Gb Ethernet, InfiniBand*, IEEE 1394b, USB 2.0, etc……

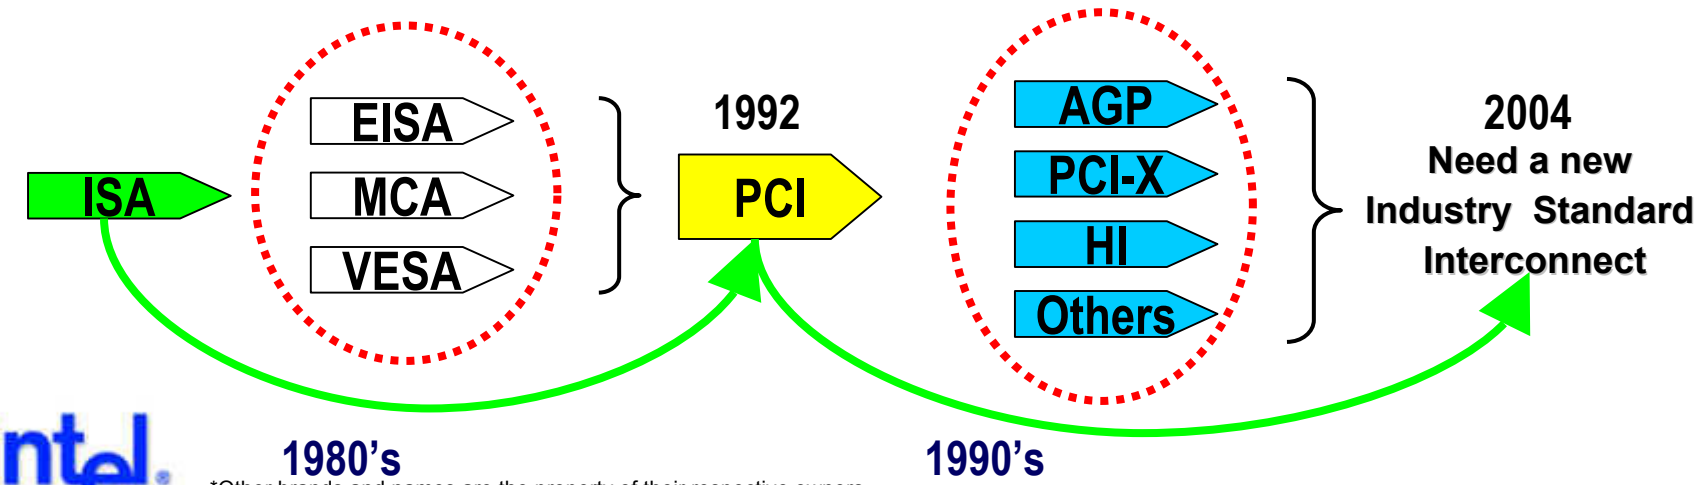

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

PCI Express* Technology Overview

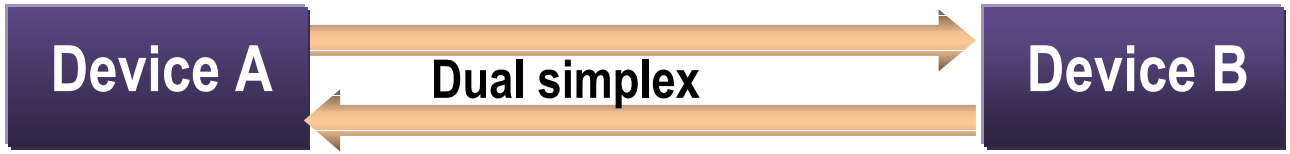

- A PCI Express* "lane" is four wires
  - One differential pair for transmit and another pair for receive
  - Signaling is at 2.5 GHz with 8b/10b encoding
- Connectors are defined for x1, x4, x8, x16 lanes providing an opportunity to scale bandwidth

Lanes	Bandwidth (peak)
x1	500 MB/s
x4	2 GB/s
x8	4 GB/s
x16	8 GB/s

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Infiniband* on PCI Express™

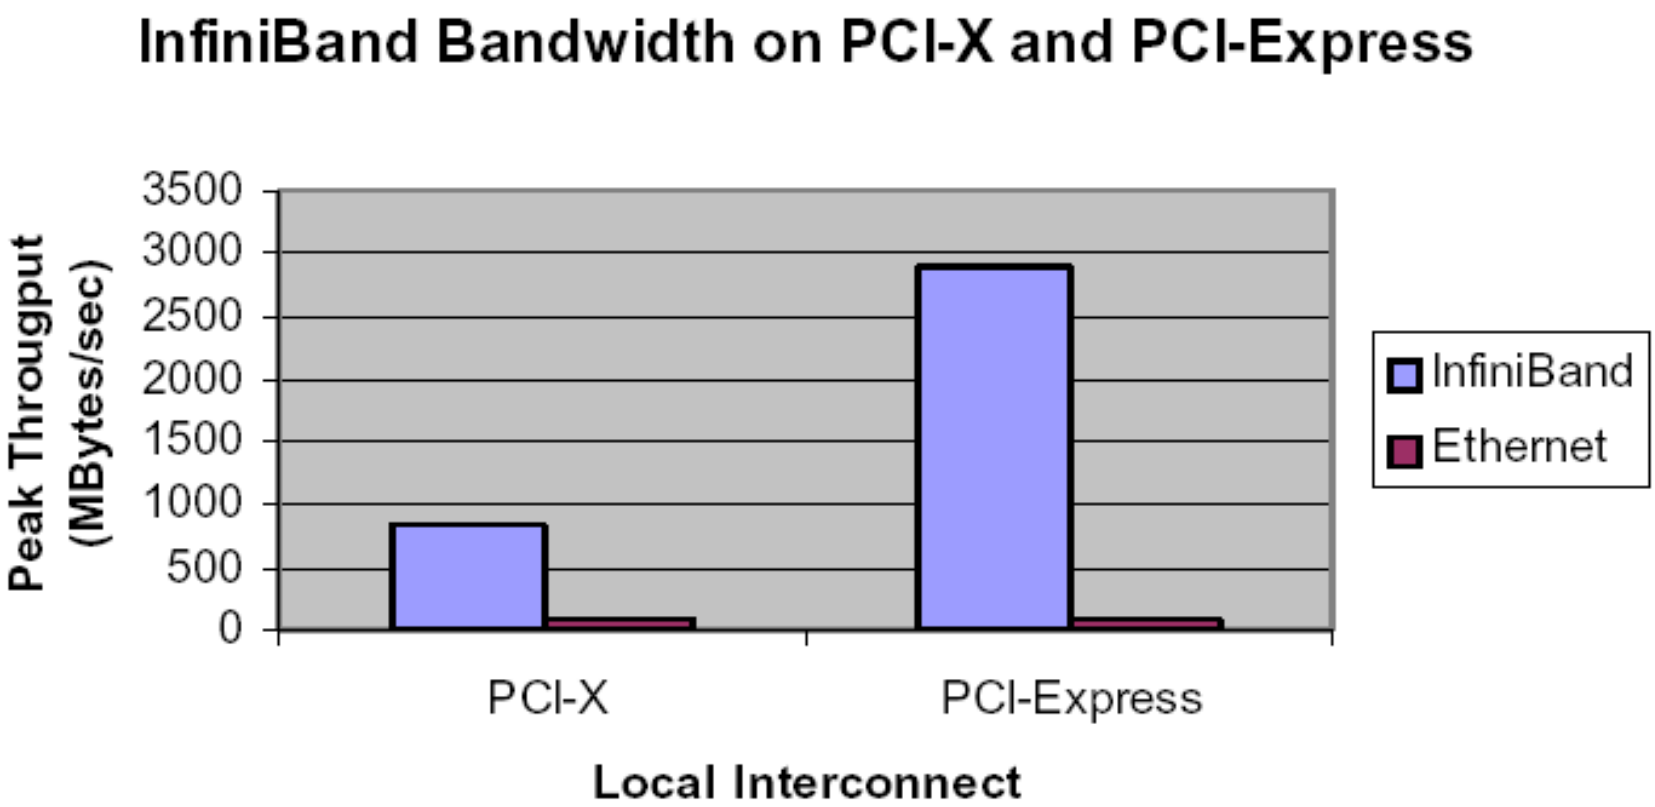

Source: Mellanox* http://www.mellanox.com/technology/shared/InfiniHost_Architecture_WP_150.pdf

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Silvervale** Technology

- New virtualization (VMX) capabilities in processors work with Virtual Machine Monitor (VMM) software to enable multiple OS instances to robustly run on virtual copies of a real complete hardware system (e.g., an IA-based platform)
  - VMX = virtual machine extensions

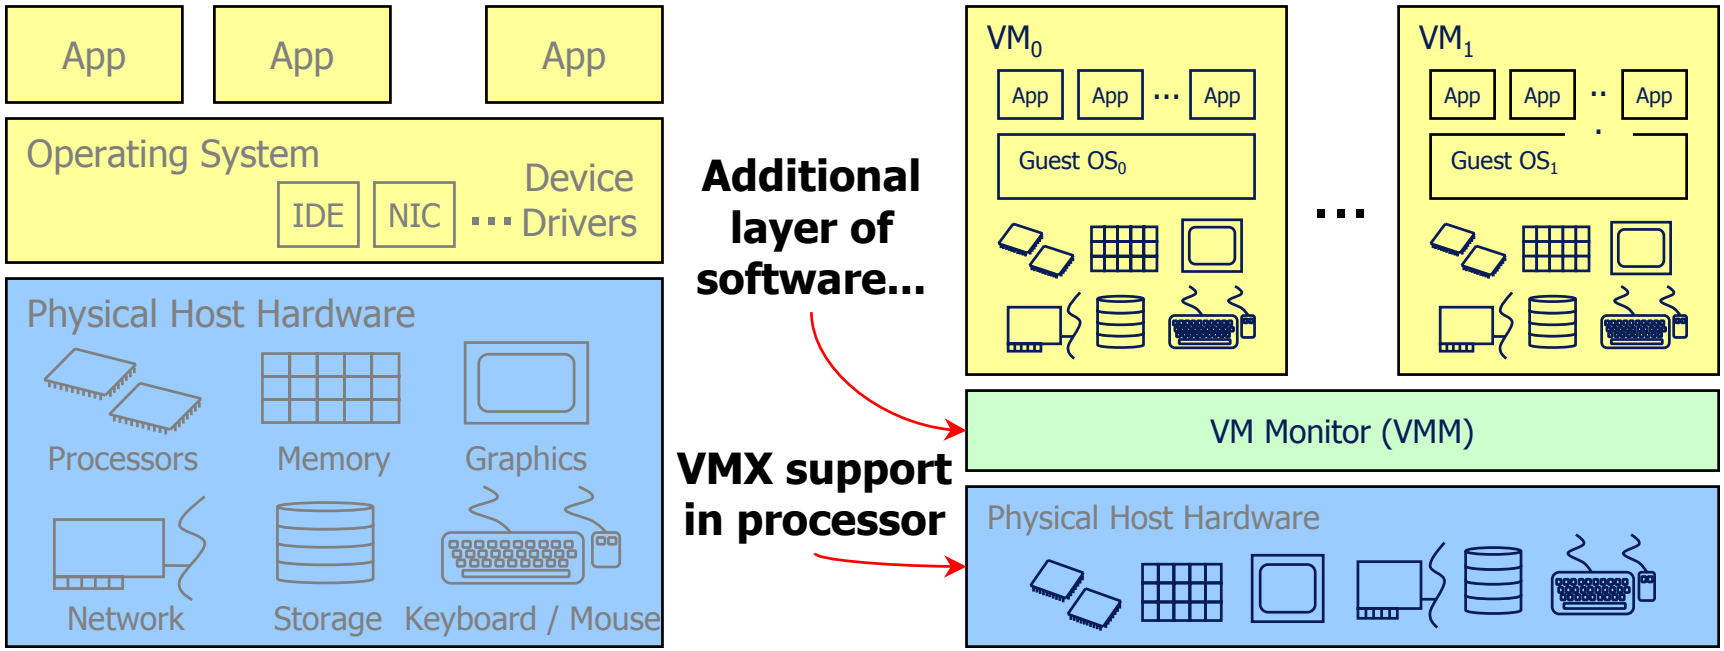

Without VMs: Single OS owns all hardware resources

With VMs: Multiple OSes share hardware resources through virtual partitions

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

**codename

Itanium® Architecture for HPC

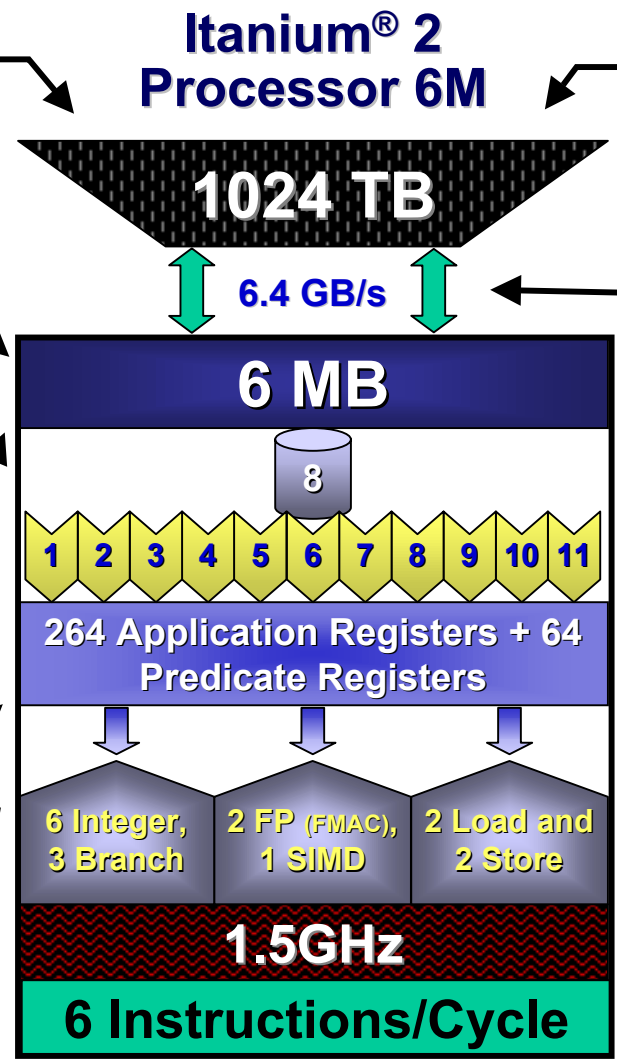

50-bit physical, 64-bit virtual addressing provides vast, flat address space to keep data close to processor

Large page size - 4GB page allows addressing 512GB w/o TLB miss

Large, low latency, on-die cache accelerates data throughput to core at a rate of 48GB/s

Advanced register model handles large, complex workloads and predicate registers accelerates loop unrolling, parallel execution

Compiler includes high precision math libraries for complex scientific calculations

High memory bandwidth for more calculations / second

Multiple issue ports / execution units avoids reservation stalls increasing parallel execution of up to 6 instructions per cycle

Dual 80-bit FMACs with 128-bit support for efficient single & double precision multiply-add

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Intel® Itanium® 2 Processor Family

Multi-Processor (MP) Capable — *Leading Performance*

2004	2005	2006	Next Generation
Itanium® 2 Processor (Madison** 9M) 1.6GHz, 9MB 1.6GHz, 6MB 1.5GHz, 4MB	Montecito** Dual Core, Larger Caches, 90nm Technology	Montvale** Dual Core	Tukwila** Multi Core, Developed with ex-Alpha team, Common Platform

Dual Processor (DP) Capable — *Lowest $ / FLOP*

2004	2005	2006	Next Generation
Itanium® 2 Processor 1.6GHz, 3MB 1.4GHz, 3MB	Millington** DP	Montvale** DP	Dimona** DP

Dual Processor (DP) Capable — *Lower Power*

2004	2005	2006	Next Generation
LV Itanium® 2 Processor 1.3GHz, 3M	LV Millington** DP, Low Voltage	LV Montvale** DP, Low Voltage	LV Dimona** DP, Low Voltage

□Low Voltage Intel® Itanium® 2 processor
** Codename

All products, dates and information are preliminary and subject to change without notice

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Montecito** Summary

- Intel's first dual core Itanium processor
- 90 nm process technology
- Intel's first processor with > 1 billion transistors
- 24MB L3 Cache
- Multithreading
- Lower power
- Socket compatible with previous Itanium 2-based systems
  - 400MHz FSB version compatible with previous Itanium 2-based systems
  - 667MHz FSB version for enabled chipsets
- Target 2005 platform release

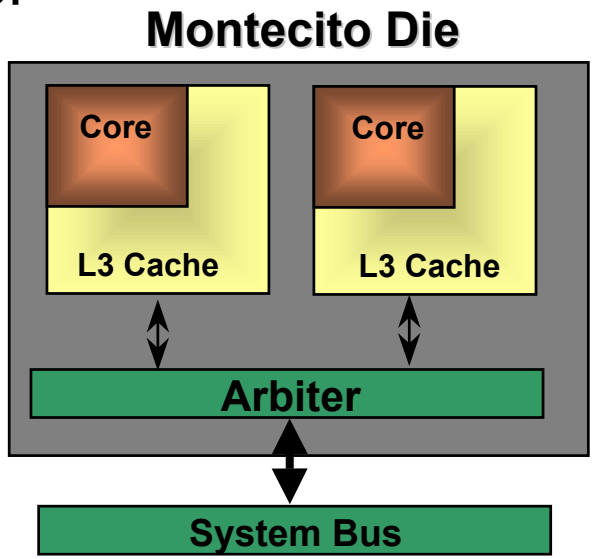

**codename

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Itanium® Architecture
Optimized for Multi-Core

- Parallel execution leadership: only Intel has all 3:
  - Multi cores on same die
  - Multi threads on same core
  - Explicit Parallelism in each core
- EPIC*: inherent advantages for multi-core, multi-thread
  - Architecture: Parallelism + many registers to keep data on-chip
  - Core size: Smaller than IA-32, up to 2x more cores per die on than on IA-32

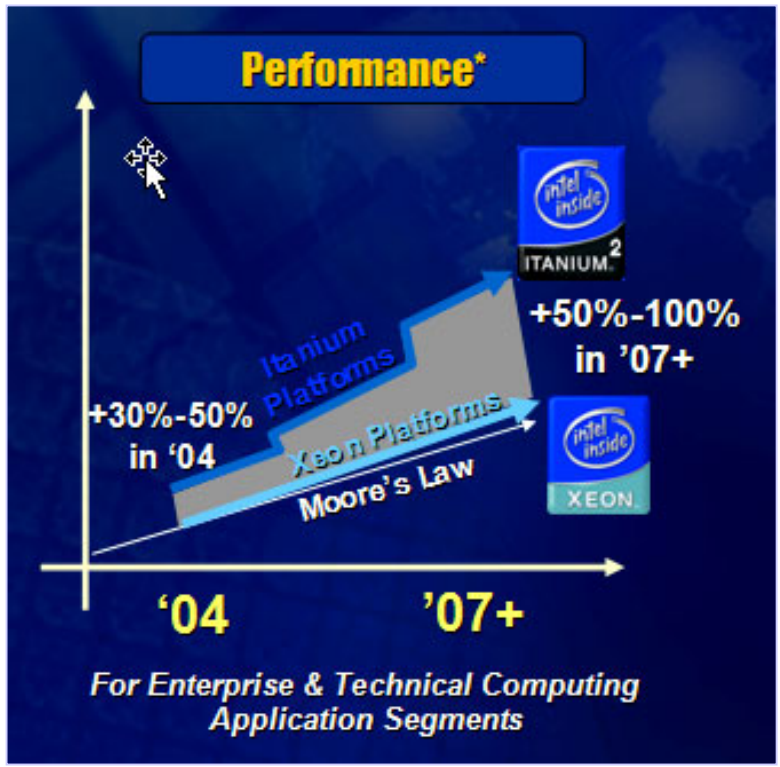

All features and dates specified are targets provided for planning purposes only and are subject to change without notice.

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Common Platform & Infrastructure

- Today: Itanium® Processor exceeds RISC performance & price / perf
- Today: Itanium® platform delivering superior price / performance vs Intel® Xeon™ Processor on transaction processing
  - 30% more transactions at 10% incremental cost of hardware platform/ OS / database***
- '07: Itanium® platform cost reduced to parity with Intel® Xeon™ processor-based platforms
  - Common platform components to lead to common platform infrastructure over time

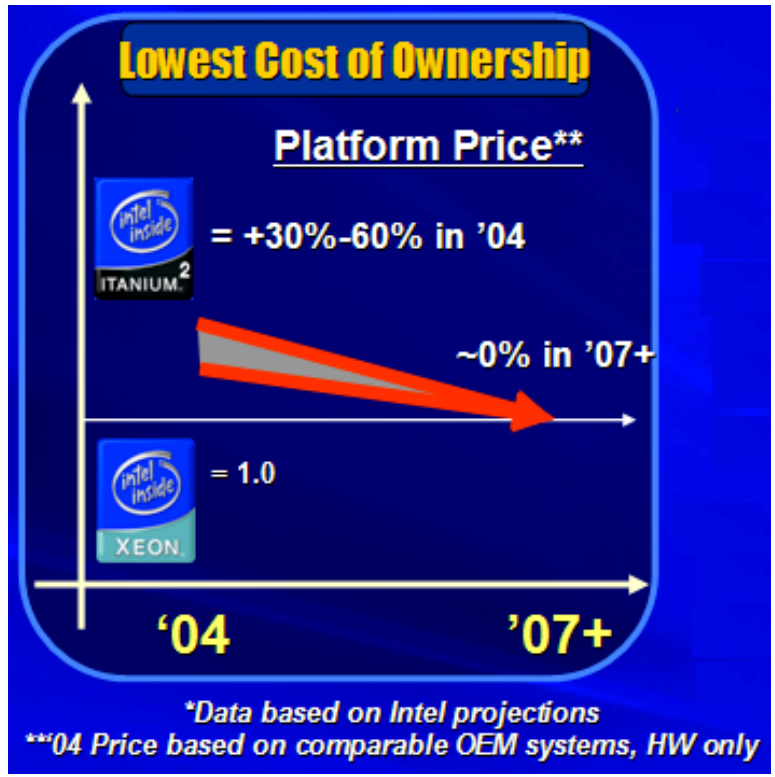

All features and dates specified are targets provided for planning purposes only and are subject to change without notice.

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.

Technology Leadership

www.intel.com

*Other brands and names are the property of their respective owners
© Copyright 2002-2003 Intel Corporation. All Rights Reserved.